

07-03-00

A

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

CERTIFICATE OF EXPRESS MAILING

I hereby certify that this paper and the documents and/or fees referred to as attached therein are being deposited with the United States Postal Service on June 30, 2000 in an envelope as "Express Mail Post Office to Addressee" service under 37 CFR §1.10, Mailing Label Number **EL269879920US**, addressed to the Assistant Commissioner for Patents, Washington, DC 20231.

Attorney Docket No.: **CISCP146**

First Named Inventor: **Martin Cieslak**



UTILITY PATENT APPLICATION TRANSMITTAL (37 CFR. § 1.53(b))
(Regular application claiming priority of a provisional)

Assistant Commissioner for Patents
Box Patent Application
Washington, DC 20231

☐ Duplicate for
fee processing

Sir: This is a request for filing a patent application under 37 CFR. § 1.53(b) in the name of inventors:

Martin Cieslak and Alex Tweedly

For: **METHOD AND APPARATUS FOR REDIRECTING NETWORK TRAFFIC**

Priority of U.S. Provisional Application No. 60/168,862 filed on December 2, 1999 is claimed under 35 U.S.C. § 119(e).

Application Elements:

- ☒ 43 Pages of Specification, Claims and Abstract
☒ 6 Sheets of **formal** Drawings
☐ Combined Declaration and Power of Attorney
☐ Separate Declaration

Accompanying Application Parts:

- ☐ Assignment and Assignment Recordation Cover Sheet (recording fee of \$40.00 enclosed)
☐ Separate Power of Attorney by Assignee
☐ 37 CFR 3.73(b) Statement by Assignee
☐ Information Disclosure Statement with Form PTO-1449
☐ Copies of IDS Citations
☐ Preliminary Amendment
☒ Return Receipt Postcard
☐ Small Entity Statement(s)
☐ Other:

Claim For Foreign Priority

☐ Priority of _____ Application No. _____ filed on _____ is claimed under 35 U.S.C. § 119.

☐ The certified copy is enclosed.

☐ The certified copy will follow.

☐ Amend the specification by inserting before the first line the sentence:

"This application claims the benefit of U.S. Provisional Application No. _____ filed on _____, the disclosure of which is incorporated herein by reference."

Fee Calculation (37 CFR § 1.16)

	(Col. 1) NO. FILED	(Col. 2) NO. EXTRA	SMALL ENTITY RATE FEE	OR	LARGE ENTITY RATE FEE
BASIC FEE			\$345 \$	OR	\$690 \$690
TOTAL CLAIMS	79 -20 = 59		x 9 = \$	OR	59 x 18 = \$1,062
INDEP CLAIMS	8 -03 = 5		x 39 = \$	OR	5 x 78 = \$390
<input checked="" type="checkbox"/> Multiple Dependent Claim Presented			\$130 = \$	OR	\$260 = \$
*If the difference in Col. 1 is less than zero, enter "0" in Col. 2.			Total \$	OR	Total \$2,142

Please defer filing fees at this time.

General Authorization for Petition for Extension of Time (37 CFR §1.136)

☒ Applicants hereby make and generally authorize any Petitions for Extensions of Time as may be needed for any subsequent filings. The Commissioner is also authorized to charge any extension fees under 37 CFR §1.17 as may be needed to Deposit Account No. 50-0388 (Order No. CISCPI46).

☒ Please send correspondence to the following address:

Customer Number 022434
BEYER WEAVER & THOMAS, LLP
P.O. Box 130
Mountain View, CA 94042-0130
Telephone (510) 843-6200
Fax (510) 843-6203



Date: 6/30/06

Mary R. Olynick
Mary R. Olynick
Registration No. 42,963

PATENT APPLICATION

METHOD AND APPARATUS
FOR REDIRECTING NETWORK TRAFFIC

Inventor(s): Martin Cieslak
1074 Canyon Creek Terrace,
Fremont, CA 94536
Citizen of USA

Alex Tweedly

Assignee: Cisco Technology, Inc.

BEYER WEAVER & THOMAS, LLP
P.O. Box 130
Mountain View, CA 94042-0130
(510) 843-6200

METHOD AND APPARATUS FOR REDIRECTING NETWORK TRAFFIC

Martin Cieslak and Alex Tweedly

CROSS REFERENCE TO RELATED APPLICATIONS

This application claims priority of U.S. provisional patent application, having Application No. 60/168,862 filed December 02, 1999 by Cieslak et al., which application is incorporated herein by reference in its entirety for all purposes.

FIELD OF THE INVENTION

The present invention relates to transmission of data in a network environment. More specifically, the present invention relates to methods and apparatus for transparently and selectively redirecting network traffic to a selected traffic handling system, such as a cache system.

BACKGROUND OF THE INVENTION

Generally speaking, when a client platform communicates with some remote server, whether via the Internet or an intranet, it crafts a data packet which defines a TCP connection between the two hosts, i.e., the client platform and the destination server. More specifically, the data packet has headers which include the destination IP address, the

destination port, the source IP address, the source port, and the protocol type. The destination IP address might be the address of a well known World Wide Web (WWW) search engine such as, for example, Yahoo, in which case, the protocol would be TCP and the destination port would be port 80, a well known port for http and the WWW. The source IP address would, of course, be the IP address for the client platform and the source port would be one of the TCP ports selected by the client. These five pieces of information define the TCP connection.

Given the increase of traffic on the World Wide Web and the growing bandwidth demands of ever more sophisticated multimedia content, there has been constant pressure to find more efficient ways to service data requests than opening direct TCP connections between a requesting client and the primary repository for the desired data. Interestingly, one technique for increasing the efficiency with which data requests are serviced came about as the result of the development of network firewalls in response to security concerns. In the early development of such security measures, proxy servers were employed as firewalls to protect networks and their client machines from corruption by undesirable content and unauthorized access from the outside world. Proxy servers were originally based on Unix machines because that was the prevalent technology at the time. This model was generalized with the advent of SOCKS which was essentially a daemon on a Unix machine. Software on a client platform on the network protected by the firewall was specially configured to communicate with the resident demon which then made the connection to a destination platform at the client's request. The demon then passed information back and forth between the client and destination platforms acting as an intermediary or "proxy."

Not only did this model provide the desired protection for the client's network, it gave the entire network the IP address of the proxy server, therefore simplifying the problem

of addressing of data packets to an increasing number of users. Moreover, because of the storage capability of the proxy server, information retrieved from remote servers could be stored rather than simply passed through to the requesting platform. This storage capability was quickly recognized as a means by which access to the World Wide Web could be accelerated. That is, by storing frequently requested data, subsequent requests for the same data could be serviced without having to retrieve the requested data from its original remote source. Currently, most Internet service providers (ISPs) accelerate access to their web sites using proxy servers.

Unfortunately, interaction with such proxy servers is not transparent, requiring each end user to select the appropriate proxy configuration in his or her browser to allow the browser to communicate with the proxy server. For the large ISPs with millions of customers there is significant overhead associated with handling tech support calls from customers who have no idea what a proxy configuration is. Additional overhead is associated with the fact that different proxy configurations must be provided for different customer operating systems. The considerable economic expense represented by this overhead offsets the benefits derived from providing accelerated access to the World Wide Web. Another problem arises as the number of WWW users increases. That is, as the number of customers for each ISP increases, the number of proxy servers required to service the growing customer base also increases. This, in turn, presents the problem of allocating packet traffic among multiple proxy servers.

Network caching represents an improvement over the proxy server model. Network caching is transparent to end users, high performance, and fault tolerant. By altering the operating system code of an existing router, the router is enabled to redirect data traffic of a particular protocol intended for a specified port, e.g., TCP with port 80, to one or more

5 caching engines connected to the router via an interface having sufficient bandwidth such as, for example, a 100baseT interface. If there are multiple caching engines connected to the cache-enabled router, the router selects from among the available caching engines for a particular request based on a simple algorithm according to which a particular group of addresses is associated with each caching engine.

10 The caching engine to which the request is re-routed "spoofs" the requested destination platform and accepts the request on its behalf via a standard TCP connection established by the cache-enable router. If the requested information is already stored in the caching engine, i.e., a cache "hit" occurs, it is transmitted to the requesting platform with a header indicating its source as the destination platform. If the requested information is not in the caching engine, i.e., a cache "miss" occurs, the caching engine opens a direct TCP connection with the destination platform, downloads the information, stores it for future use, and transmits it to the requesting platform. All of this is transparent to the user at the requesting platform which operates exactly as if it were communicating with the destination platform. Thus, the need for configuring the requesting platform to suit a particular proxy 15 configuration is eliminated along with the associated overhead. Moreover, traffic may be easily allocated among as many caching engines as become necessary. Thus, content caching provides a way to compensate for the bandwidth limitations discussed above. Specific examples of such a network caching technique are embodied in the Web Cache Coordination Protocol (WCCP) by Cisco Systems, which examples are further described in 20 U.S. Patent Application No. 08/946,867 for METHOD AND APPARATUS FOR FACILITATING NETWORK DATA TRANSMISSIONS filed October 8, 1997. This patent application is herein incorporated by reference in its entirety for all purposes.

Although the traffic redirection mechanisms described above in relation to cache engines works well, these mechanisms provide limited traffic redirection capabilities. For example, only packets that are destined for port 80 are redirected. Additionally, all of the traffic destined for port 80 is automatically redirected without intelligently determining whether to redirect portions of such traffic. Once the cache system receives the redirected traffic, there are no mechanisms for refusing the redirected traffic. The cache system may be unable to handle such redirected traffic and accordingly redirected traffic may be halted because the destination cache system is unavailable.

Accordingly, there is a need for improved mechanisms for intelligently redirecting traffic (*e.g.*, based on traffic characteristics). More specifically, there is a need for capabilities that allow a cache system (or other type of computer system) to refuse a redirected packet. Finally, there is also a need for mechanisms for processing and/or monitoring the redirected traffic.

SUMMARY OF THE INVENTION

Accordingly, the present invention provides an apparatus and method for intelligently and transparently redirecting and/or processing traffic. In one embodiment, a method of facilitating redirection of traffic sent from a first processing device (e.g., a client) to a second
5 processing device (e.g., a server) is disclosed. At a third processing device (e.g., a router) associated with a plurality of traffic handling systems (e.g., cache systems), traffic information is received from at least one of the associated traffic handling systems. A determination as to how to redirect data received by the third processing device to a selected traffic handling system may be based on the traffic information received from at least one of
10 the associated traffic handling systems. That is, each traffic handling system may specify how traffic is to be redirected to such traffic handling system. For example, a traffic handling system may select traffic destined for a particular port number.

In a preferred embodiment, the determination of redirecting data is accomplished by communicating the traffic information to at least a designated one of the associated traffic
15 handling systems and receiving traffic redirection information from the designated traffic handling system at the third processing device. For example, the third processing device compiles traffic information received from each traffic handling system and sends the compiled traffic information to a designated traffic handling system. The designated traffic handling system then generates the traffic redirection information, which is based on the
20 communicated traffic information. In one aspect, a data structure (e.g., a table that lists traffic information from each traffic handling system) may be built and updated based on the received traffic information. The traffic information is communicated to the designated traffic handling system within the data structure. In another aspect, when a packet is received in the third processing device from the first processing device destined for the

second processing device, the packet may then be redirected to a selected one of the traffic handling systems based on the traffic redirection information.

In another method implementation, traffic information is received at a designated traffic handling system associated with a plurality of traffic handling systems and one or more routers from at least one of the routers. The traffic information specifies which data should be redirected to which associated traffic handling systems. Traffic portions are allocated to each traffic handling system based on the received traffic information. Preferably, the traffic portions are allocated by building or updating a data structure based on the traffic information received within the third processing device. The data structure specifies how to direct traffic received in the third processing device to selected traffic handling systems associated with the third processing device.

In another implementation, when a traffic handling system receives redirected traffic, the receiving traffic handling system may transmit particular data portions back to the redirecting router so that the data is reinserted into the original traffic flow. This may be done in a manner transparent to the source or destination of the packets. An example of a case in which packets would be reinserted in the original flow might be where the cache recognizes the source and destination pairs identified by the packets as corresponding to a connection requiring IP-based authentication. Another example would be where the cache is overloaded and is currently unable to handle all of the redirected traffic.

In another embodiment, the invention pertains to a computer system operable to facilitate redirection of traffic sent from a first processing device to a second processing device. The computer system includes a memory and a processor coupled to the memory. The memory and the processor are adapted to provide at least some of the above described

method operations. In yet a further embodiment, the invention pertains to a computer program product for facilitating redirection of traffic sent from a first processing device to a second processing device. The computer program product has at least one computer readable medium and a computer program instructions stored within the at least one computer readable product configured to cause a processing device to perform at least some of the above described method operations.

These and other features and advantages of the present invention will be presented in more detail in the following specification of the invention and the accompanying figures which illustrate by way of example the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a network diagram illustrating a plurality of cache systems according to a specific embodiment of the present invention;

FIG. 2 is a flowchart illustrating a procedure for traffic redirection in accordance with one embodiment of the present invention.

FIG. 3 is a flowchart illustrating the operation of FIG. 2 for sending Traffic Information from a cache system to one or more routers within the farm in accordance with one embodiment of the present invention.

FIG. 4 is a flowchart illustrating both of the operations of FIG. 2 for periodically updating a Router View and sending it from a router to one or more cache systems within a cluster in accordance with one embodiment of the present invention.

FIG. 5 is a flowchart illustrating operation of FIG. 4 for validating the `HERE_I_AM` message received from a cache system in accordance with one embodiment of the present invention.

FIG. 6 is a diagrammatic representation of a router in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

Reference will now be made in detail to a specific embodiment of the invention. An example of this embodiment is illustrated in the accompanying drawings. While the invention will be described in conjunction with this specific embodiment, it will be understood that it is not intended to limit the invention to one embodiment. On the contrary, it is intended to cover alternatives, modifications, and equivalents as may be included within the spirit and scope of the invention as defined by the appended claims. In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. The present invention may be practiced without some or all of these specific details. In other instances, well known process operations have not been described in detail in order not to unnecessarily obscure the present invention.

FIG. 1 is a simplified network diagram which will be used in conjunction with the flowcharts of FIGS. 2 and 3 to describe a specific embodiment of the present invention. According to this embodiment, a plurality of client machines 102 which are resident on local area networks (LAN) 104a and 104b communicate via router 106/router 114 and wide area network (WAN) 108, e.g., the internet, with server 110. Of course, some or all of the clients 102 may communicate with the router 106 or 114 through various other configurations, rather than through a LAN. For example, a client may be coupled directly to the router 106 or there may be one or more intermediate routers between a client 102 and the router 106.

As discussed above, the router 106 may direct certain traffic, e.g., destined for port 80, to a cache system, such as 112a, which is configured to "spoof" server 110. If there are multiple caches connected to the cache-enabled router, the router selects from among the available caches for a particular request based on any suitable factor as described below with

reference to FIGS. 2 through 5. For example, a first set of destination IP addresses may be assigned to cache system 112a; a second set of IP addresses to cache system 112b; a third set of IP addresses to cache system 112c; and a fourth set of IP addresses to cache system 112d.

Before sending the packet to one of its associated cache systems, the cache-enabled router 106 encapsulates the packet for transmission to the selected cache system by adding another TCP/IP header which designates the router as the source of the packet and the cache system as the destination. That is, the router encapsulates the packet for transmission to a cache system which might be several "hops" away. So, for example, router 106 might encapsulate the packet for transmission to cache system 112d which is connected to router 106 via router 114. Thus, not only may multiple cache systems be associated with a particular router, but multiple routers may be supported by an individual cache system or a group of cache systems. This allows a tremendous amount of flexibility in where the cache system and router need to be in relation to each other.

The selected cache system 112a may respond to a request from a client 102 to obtain objects from destination platform 110. The cache system 112a may retrieve objects from destination platform 110 to then present to one of the clients or may retrieve objects from its own cache (which objects were previously retrieved from the destination platform 110).

It will be understood that the network cache systems described herein may employ any of a variety of existing file systems and remain within the scope of the invention. For example, the invention may be implemented using a Unix general purpose file system or the equivalent. A particular embodiment of the invention employs the file system described in commonly assigned, U.S. Patent No. 5,950,205 for DATA TRANSMISSION OVER THE

INTERNET USING A CACHE MEMORY FILE SYSTEM issued on September 7, 1999,
the entire specification of which is incorporated herein by reference for all purposes.

In the illustrated embodiment, cache systems 112a, 112b, 112c, and 112d form a
cache cluster or farm 120. Traffic is typically allocated to each cache system within the
same cache cluster. In general terms, the present invention provides mechanisms for
allocating traffic to each cache system within the same cluster. For example, traffic is
allocated based on IP destination address or source address. That is, each cache system may
be assigned to handle requests for objects from a particular set of destination or source
addresses. After the traffic is allocated to each cache system, mechanisms are also provided
for a cache system to send redirected traffic back to an associated router within the cluster,
which then sends the data to its specified destination address. The cache system and/or the
router may also perform any suitable number and type of data processing on such resent
data, *e.g.*, prior to reinserting the data.

FIG. 2 is a flowchart illustrating a procedure 200 for traffic redirection in accordance
with one embodiment of the present invention. Initially, each cache system within a farm or
cluster periodically sends Traffic Information to one or more routers in operation 202. In
general terms, the Traffic Information specifies what data parameters will be utilized to
determine redirection for such cache system. For example, data having a particular hashed
source IP address may be sent to a particular cache system. By way of another example,
data having a particular port source number may be redirected to a particular cache system.

The present invention provides a significant amount of flexibility in the kinds of
services offered by the cache systems or traffic handling systems. A same router may
support multiple services (*e.g.*, specified by multiple Traffic Information). For example,

router 106 may support web caching, as well as DNS caching. A same cache system may also support different services. That is, each cache system may perform any suitable traffic handling procedures (instead of or in addition to caching). By way of examples, a particular cache system may perform one or more of the following traffic handling services:

5 monitoring traffic type and/or quantity, implementing security measures, performing standard server services, caching (e.g., web, news, ftp), and forwarding and/or filtering of received traffic. This list is not meant to be exhaustive, and therefore is not meant to limit the scope of the invention. Also, each cache system may receive different types of traffic. By way of a final example, a cache may easily be set up to perform a reverse caching

10 procedure.

Each router then builds a Router View based on the Traffic Information received from each cache system in operation 204. Thus, a particular Router View of a particular router specifies and identifies the cache systems and their particular Traffic Information that are associated with such router. Each router also sends the Router View to each cache in

15 response to the Traffic Information in operation 204. Additionally, each router tracks the member cache system's availability and updates the Router View as needed in operation 206. Accordingly, each cache system can know which cache systems and Traffic Information are associated with which routers within its particular cache cluster or farm.

A designated cache system may then direct the routers of the cluster in building the

20 router's Redirection Hash Table based on the Router View received from the routers in operation 208. The Redirection Hash Table generally specifies which traffic is to be redirected to which cache system within a particular cache cluster. Each router may then redirect traffic based on its Redirection Hash Table in operation 210, and the traffic redirection procedure 200 repeats.

FIG. 3 is a flowchart illustrating the operation 202 of FIG. 2 for sending Traffic Information from a cache system to one or more routers within the farm in accordance with one embodiment of the present invention. Initially, security options are set in operation 302. Security options may include any suitable parameters for specifying security options. In a specific embodiment, the security options include a no security option and a MD5 password security option. The latter option requires that each router and cache system wishing to join a cluster be configured with the cluster password. Each WCCP protocol packet sent by a router or cache system for that cluster will contain in its security component the MD5 checksum of the WCCP protocol message and the cluster password. Each cache system and router in the cluster then authenticates the security component in a received protocol packet immediately after validating the WCCP message header. Packets failing authentication are discarded.

Service options are then set in operation 304. The service options may specify which packets are to be redirected to the particular cache system using any suitable one or more parameter(s). In the illustrated example, the service options includes a number of flags or fields for selecting which traffic to redirect to the cache system. These flags or fields may, for example, specify specific portions of the received packet that are to be hashed to create a first index to the Redirection Hash Table. In one specific implementation, any suitable combination of the following packet fields may be XOR'd together and then hashed to produce the first index: source IP address, destination IP address, source port, and destination port. The flags or fields may also specify portions of the received packet that are to be hashed to create a second index that may be used in conjunction with the first index to distribute a single bucket among a plurality of cache systems. Additionally, particular packets having specified fields (e.g., source port 8) may also be redirected to the cache

system. Several service option embodiments are further described in the above referenced parent U.S. provisional Application. In sum, the service options provide a high degree of flexibility and allow each cache system to specify what traffic is to be redirected to it.

The cache ID, which uniquely identifies the cache system, is also set in operation 306. The Cache's View information is also updated or set in operation 308. The Cache View is analogous to the Router View, but includes information (*e.g.*, cache system identifications and associated service options) specified in previously received Router Views from routers within the cache system's cluster. Initially, the Cache View may be empty or merely contain the cache's own Traffic Information.

A predetermined time count is then set in operation 310. In a specific embodiment, the predetermined time count is set to 10 seconds. The predetermined time count is then decremented in operation 312. It is then determined whether the time count has expired in operation 314. Alternatively, when the counter reaches zero, an interrupt routine may be initiated. After the time count has expired, the cache system then sends a "HERE_I_AM" message, which includes Traffic Information, to its associated routers in operation 316. The procedure then returns to operation 310 where the predetermined time count is reset. Thus, a HERE_I_AM message is sent every predetermined time count.

However, if it is determined that the predetermined time count has not expired in operation 314, it is then determined whether an I_SEE_YOU message is received in operation 318. A received I_SEE_YOU message contains the associated Router's View. If an I_SEE_YOU message is received, the cache's View information is then updated with this new received Router View information in operation 308 and the predetermined time count is decremented in operation 314. If an I_SEE_YOU message has not been received in

operation 318, the cache view update operation 308 is skipped and the predetermined time count is decremented in operation 312. The time count continues to be decremented until it expires so that a HERE_I_AM message may be sent.

Each cache system within a particular cluster may implement the procedure of FIG. 3

5 to send HERE_I_AM messages to the routers associated with the particular cluster and receive I_SEE_YOU messages in response to the HERE_I_AM messages.

FIG. 4 is a flowchart illustrating both of the operations 204 and 206 of FIG. 2 for periodically updating a Router View and sending it from a router to one or more cache systems within a cluster in accordance with one embodiment of the present invention. In general terms, each router builds and updates a Router View based on Traffic Information received from associated (or member) cache systems and periodically tracks each member cache system's availability. The updated Router View is then sent to each member cache system.

Initially, a Receive ID variable is cleared to zero in operation 402. The Receive ID variable is generally used to verify the authenticity of a HERE_I_AM message received from a cache system in response to an I_SEE_YOU message from the router by matching the Receive ID variables of the sent I_SEE_YOU message and a subsequently received HERE_I_AM message. A Change Number variable is also cleared in operation 404. The Change Number variable generally indicates the number of times that the membership has changed for a particular router. Thus, this variable is incremented each time the membership changes (e.g., a Cache System come online and sends a HERE_I_AM message). A Query Timeout variable and a Connection Timeout variable are set for each cache system member

in operation 406. These timeout variables then begin to count down and are used to determine whether a cache system has timed out.

It is then determined whether a `HERE_I_AM` message is received from a cache system in operation 407. If a message is not received, it is then determined whether the Query Timeout variable has expired in operation 412. If it has not expired, the router continues to check whether a `HERE_I_AM` message is received from a cache system in operation 407. However, if the Query Timeout variable has expired, it is then determined whether the Connection Timeout variable has expired in operation 414. If the Connection Timeout variable has expired, the expired cache is then removed from the Router's View in operation 420. The Change No. Variable is then incremented in operation 411. The Change No. Variable generally corresponds to the number of times that the number of available/valid cache systems has changed (positive and negative changes). The procedure then again waits for a next `HERE_I_AM` message in operation 407.

If the Connection Timeout variable has not expired and the Query Timeout variable has expired, however, a query is then sent to the expired cache system that has not sent a `HERE_I_AM` message in a timely manner in operation 416. For example, an `I_SEE_YOU` message may be sent. The procedure then again waits for a next `HERE_I_AM` message in operation 407.

Turning back to operation 407, when it is determined that a `HERE_I_AM` message is received from a cache system, it is then determined whether the `HERE_I_AM` message is valid in operation 408. If the message is not valid, the invalidated cache is removed from the router's View in operation 420 and the Change No. Variable is incremented in operation 411. However, if the `HERE_I_AM` message is valid, the router's View is updated based on

the HERE_I_AM message in operation 410. The procedure then again waits for a next HERE_I_AM message in operation 407.

FIG. 5 is a flowchart illustrating the operation 408 of FIG. 4 for validating the HERE_I_AM message received from a cache system in accordance with one embodiment of the present invention. Initially, the Receive ID is incremented in operation 502. The router then sends an I_SEE_YOU message with the incremented Receive ID and Change number to the Cache System in operation 504. Timeout variables for this cache system, which is the destination of the I_SEE_YOU message, are then reset in operation 506. It is then determined whether the received HERE_I_AM message is from the same cache system from which the previous HERE_I_AM message was sent (and to which the I_SEE_YOU was sent) in operation 508. If the HERE_I_AM message is not from the same cache system, it is determined that the HERE_I_AM message is not valid in operation 518 and the validity procedure ends.

If the HERE_I_AM message is from the same cache system, it is then determined whether the service indicated within the HERE_I_AM message from the cache system is valid in operation 509. For example, it is determined whether the specified service is supported by the particular router and/or other cluster routers. If the service is valid, it is then determined whether there is a match between the Receive IDs within the currently received HERE_I_AM message and the previously sent I_SEE_YOU message in operation 510. If there is a match, it is then determined whether this is the first valid match in operation 511. If this is the first valid match, the Change No. Variable is incremented in operation 512 since the number of available cache systems has changed (*i.e.*, increased by one). It is then determined that the HERE_I_AM message is valid in operation 514.

If the service is not valid in operation 509 or the Receive ID's fail to match in operation 510, it is then determined that the member is no longer valid in operation 518. If this was not the first valid match in operation 511, the operation 512 for incrementing the Change No. is skipped and the HERE_I_AM is determined to be valid in operation 514.

When the HERE_I_AM message is found to be valid, the Router View is then updated in operation 410 of FIG. 4. When the HERE_I_AM message is found to be invalid, however, the particular cache system is removed from the Router View in operation 420 of FIG. 4. The Change No. Variable is also incremented in operation 411 when the cache system is removed from the Router View since the number of available cache systems has changed (*i.e.*, decreased by one).

The above-described mechanisms facilitate traffic redirection by one or more routers to a plurality of cache systems within a particular cluster or farm. Each cache system within the cluster may specify what data should be redirected (*e.g.*, via the above described service options) to the particular cache system. Each router may then receive specifications for redirection for each cache system within the cluster. Each router may then broadcast the specifications (*e.g.*, via the Router View) for each cache system to each cache system within the cluster so that each cache system becomes aware of the available routers and associated cache systems within the cluster and their redirection specifications.

A designated cache system may then utilize the specifications from each cache system to allocate traffic to each cache system. For example, a first cache system may receive data destined for a first port number and a second cache system may receive data destined for a second port number. The designated cache system may be selected in any suitable manner. In one specific implementation, the cache system having the lowest IP

address is the designated cache system. The designated cache system may allocate data using any suitable technique so that cache system availability and traffic redirection specifications by such cache systems is taken into account. For example, several techniques for allocating buckets are described in concurrently filed U.S. patent application, having
5 Application No. _____ (Attorney Docket No. CISCPI51) entitled METHODS AND APPARATUS FOR SLOW-STARTING A WEB CACHE SYSTEM by Tiwana et al., which application is herein incorporated by reference in its entirety. These bucket allocation techniques may be easily modified so that bucket allocation is also based on data redirection specifications from each cache system. That is, bucket allocation and assignment is also
10 based on each cache system's redirection specifications.

Once a cache system receives redirected data, it may send the data back to the router, for example. That is, the data is reinserted into the original traffic flow. This may be done in a manner transparent to the source or destination of the packets. An example of a case in which packets are reinserted in the original flow is where the cache recognizes the source
15 and/or destination addresses of the packets as corresponding to a connection requiring IP-based authentication. Another example would be where the cache is overloaded and is currently unable to handle all of the redirected traffic. Prior to reinserting the data, the data may easily be modified for any suitable reason (e.g., by the cache system or router). For example, a prohibited destination address (e.g., pornography.com) may be modified to a
20 different destination address (e.g., Disneyland.com ®). By way of another example, informative tags may be inserted into the packet for specifying any suitable identification parameter (e.g., such as the source's geographical location). Additionally, copies of packets or packet portions may be retained for various monitoring purposes (e.g., to monitor potential hacker or terrorist activities).

The cache system may send a redirected packet back to the router utilizing any suitable technique so that the router does not attempt to redirect such returned packets. However since the returned packet will retain its original source IP address, the router cannot use this to distinguish these packets from valid candidates for redirection. There are two mechanisms to prevent redirection of returned packets: Interface-based and Encapsulated-based. The Interface mechanism requires a router to be configured to inhibit redirection of packets arriving over interfaces connected to web caches. This mechanism is efficient but is topology-dependant and thus will not always be suitable. The Encapsulation mechanism requires a web-cache to send encapsulated returned packets to the router. For example, encapsulated return packets are encapsulated in a GRE packet with a Protocol Type of 0x883E and contain the original Redirect Header. The router decapsulates the packets and forwards them without attempting to redirect.

Generally, portions of the traffic redirection techniques of the present invention may be implemented on software and/or hardware. For example, they can be implemented in an operating system kernel, in a separate user process, in a library package bound into network applications, on a specially constructed machine, or on a network interface card. In a specific embodiment of this invention, a technique of the present invention is implemented in software such as an operating system or in an application running on an operating system.

A software or software/hardware hybrid traffic redirection system of this invention is preferably implemented on a general-purpose programmable machine selectively activated or reconfigured by a computer program stored in memory. The programmable machine may be in the form of a general purpose computer that is configured to implement a cache

system. The computer system generally includes a memory, processor, and network interface modules.

Such programmable machine may also be a network device designed to handle network traffic. Such network devices typically have multiple network interfaces including frame relay and ISDN interfaces, for example. Specific examples of such network devices include routers and switches. For example, the traffic redirection systems of this invention may be specially configured routers such as specially configured router models 1600, 2500, 2600, 3600, 4500, 4700, 7200, 7500, and 12000 available from Cisco Systems, Inc. of San Jose, California. A general architecture for some of these machines will appear from the description given below. In an alternative embodiment, a traffic redirection system may be implemented on one or more general-purpose network host machine such as a personal computer or workstation. Further, techniques of the invention may be at least partially implemented on a card (e.g., an interface card) for a network device or a general-purpose computing device.

Referring now to FIG. 6, a router 10 suitable for implementing portions of the present invention includes a master central processing unit (CPU) 62, interfaces 68, and a bus 15 (e.g., a PCI bus). When acting under the control of appropriate software or firmware, the CPU 62 is responsible for such router tasks as routing table computations and network management. It may also be responsible for determining the tasks described above for the client gateway router, the web server, the switch, the server gateway router, or each replica. It preferably accomplishes all these functions under the control of software including an operating system (e.g., the Internetwork Operating System (IOS®) of Cisco Systems, Inc.) and any appropriate applications software. CPU 62 may include one or more processors 63

such as a processor from the Motorola family of microprocessors or the MIPS family of

microprocessors. In an alternative embodiment, processor 63 is specially designed hardware for controlling the operations of router 10. In a specific embodiment, a memory 61 (such as non-volatile RAM and/or ROM) also forms part of CPU 62. However, there are many different ways in which memory could be coupled to the system. Memory block 61 may be used for a variety of purposes such as, for example, caching and/or storing data, programming instructions, etc.

The interfaces 68 are typically provided as interface cards (sometimes referred to as "line cards"). Generally, they control the sending and receiving of data packets over the network and sometimes support other peripherals used with the router 10. Among the interfaces that may be provided are Ethernet interfaces, frame relay interfaces, cable interfaces, DSL interfaces, token ring interfaces, and the like. In addition, various very high-speed interfaces may be provided such as fast Ethernet interfaces, Gigabit Ethernet interfaces, ATM interfaces, HSSI interfaces, POS interfaces, FDDI interfaces and the like. Generally, these interfaces may include ports appropriate for communication with the appropriate media. In some cases, they may also include an independent processor and, in some instances, volatile RAM. The independent processors may control such communications intensive tasks as packet switching, media control and management. By providing separate processors for the communications intensive tasks, these interfaces allow the master microprocessor 62 to efficiently perform routing computations, network diagnostics, security functions, etc.

Although the system shown in FIG. 6 is one specific router of the present invention, it is by no means the only router architecture on which the present invention can be implemented. For example, an architecture having a single processor that handles

communications as well as routing computations, etc. is often used. Further, other types of interfaces and media could also be used with the router.

Regardless of network device's configuration, it may employ one or more memories or memory modules (such as, for example, memory block 65) configured to store data, program instructions for the general-purpose network operations and/or bucket allocation and assignment mechanisms described herein. The program instructions may control the operation of an operating system and/or one or more applications, for example. The memory or memories may also be configured to store the above described Redirection Table, Router View, or Cache View.

Because such information and program instructions may be employed to implement the systems/methods described herein, the present invention relates to machine readable media that include program instructions, state information, etc. for performing various operations described herein. Examples of machine-readable media include, but are not limited to, magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical media such as floptical disks; and hardware devices that are specially configured to store and perform program instructions, such as read-only memory devices (ROM) and random access memory (RAM). The invention may also be embodied in a carrier wave travelling over an appropriate medium such as airwaves, optical lines, electric lines, etc. Examples of program instructions include both machine code, such as produced by a compiler, and files containing higher level code that may be executed by the computer using an interpreter.

Although the foregoing invention has been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be

practiced within the scope of the appended claims. Therefore, the described embodiments should be taken as illustrative and not restrictive, and the invention should not be limited to the details given herein but should be defined by the following claims and their full scope of equivalents.

100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000

WHAT IS CLAIMED IS:

1. A method of facilitating redirection of traffic sent from a first processing device to a second processing device, the method comprising:

at a third processing device associated with a plurality of traffic handling systems, receiving traffic information from at least a portion of the associated traffic handling systems, wherein the traffic information specifies which data should be redirected to the portion of associated traffic handling systems; and

determining how to redirect data received by the third processing device to a selected traffic handling system based on the received traffic information.

2. A method as recited in claim 1, wherein the determination of redirecting data is accomplished by:

communicating the traffic information to at least a designated one of the associated traffic handling systems; and

at the third processing device, receiving traffic redirection information from the designated traffic handling system, the traffic redirection information being based on the communicated traffic information.

3. A method as recited in claim 1, further comprising:

at the third processing device, building or updating a data structure based on the received traffic information, wherein the traffic information is communicated to the designated traffic handling system within the data structure.

4. A method as recited in claim 1, further comprising:

in the third processing device, receiving a packet from the first processing device destined for the second processing device; and

redirecting the packet to a selected one of the traffic handling systems based on the traffic redirection information.

5. A method as recited in claim 4, further comprising:

receiving the packet back after redirecting it to the selected traffic handling system;

determining that the packet is to be sent to the packet's original destination address instead of being redirected to the selected traffic handling system; and

sending the packet to its original destination.

6. A method as recited in claim 5, wherein it is determined that the packet is to be sent to the packet's original destination by determining that the packet is encapsulated and de-encapsulating the packet prior to sending the packet to its original destination.

7. A method as recited in claim 1, wherein the traffic information sent from a selected traffic handling system to the third processing device includes service options specifying which data is to be redirected to the selected traffic handling system.

8. A method as recited in claim 7, wherein the service options include a plurality of fields that are configurable to indicate that one or more fields of a packet received in the third processing device are to be used to determined redirection of packets to the selected traffic handling system.

9. A method as recited in claim 8, wherein the fields are selected from a group consisting of a source IP field, a destination IP field, a source port field, a destination port field, a source IP alternative field, a destination IP alternative field, a source port alternative field, and a destination port alternative field.

10. A method as recited in claim 9, wherein each field indicates that a corresponding field of a packet received in the third processing device is to be used to generate an index to a table identifying the plurality of associated traffic handling systems, the generated index being associated with the selected traffic handling system.

11. A method as recited in claim 10, wherein each field indicates that a hashed value of the corresponding field of the received packet is to be used to generate the index to the table identifying the plurality of associated traffic handling systems.

12. A method as recited in claim 8, wherein at least one of the fields may be set to indicate one or more port identifiers of traffic received in the third processing device.

13. A method as recited in claim 8, wherein the fields are selected from a group consisting of a port 0 field, a port 1 field, a port 2 field, a port 3 field, a port 4 field, a port 5 field, a port 6 field, and a port 7 field.

14. A method as recited in claim 13, wherein the fields includes a source/destination field to indicate whether the port identifiers of the received traffic are source ports or destination ports.

15. A method as recited in claim 6, further comprising:

in the third processing device, receiving a packet from the first processing device destined for the second processing device; and

when one or more port identifiers of the received packet matches a corresponding set field of the service options of the selected traffic handling system, redirecting the packet to the selected traffic handling system.

16. A method as recited in claim 1, wherein the traffic information sent from a selected traffic handling system to the third processing device includes security options for specifying an authentication level for messages communicated between the third processing device and the selected traffic handling system.

17. A method as recited in claim 16, wherein the security options are configurable to select no authentication for messages communicated between the third processing device and the selected traffic handling system.

18. A method as recited in claim 16, wherein the security options are configurable to require a predetermined password encoded within messages communicated between the third processing device and the selected traffic handling system.

19. A method as recited in claim 1, wherein the traffic information sent from a selected traffic handling system includes identifying information for the selected traffic handling system.

20. A method of facilitating redirection of traffic sent from a first processing device to a second processing device, the method comprising:

at a designated traffic handling system associated with a plurality of traffic handling systems and one or more routers, receiving traffic information from at least one of the routers, wherein the traffic information specifies

which data should be redirected to which associated traffic handling systems;
and

allocating traffic portions to each traffic handling system based on the
received traffic information.

21. A method as recited in claim 20, wherein the traffic portions are allocated by:
building or updating a data structure based on the received traffic
information within a third processing device that receives traffic sent from a
first processing device to a second processing device, the data structure
specifying how to direct traffic received in the third processing device to
selected traffic handling systems associated with the third processing device.

22. A method as recited in claim 20, wherein the traffic information includes
service options specifying which data is to be redirected to the selected traffic handling
system.

23. A method as recited in claim 22, wherein the service options include a
plurality of fields that are configurable to indicate that one or more fields of a packet
received in the third processing device are to be used to determined redirection of packets to
the selected traffic handling system.

24. A method as recited in claim 23, wherein the flags are selected from a group
consisting of a source IP field, a destination IP field, a source port field, a destination port
field, a source IP alternative field, a destination IP alternative field, a source port alternative
field, and a destination port alternative field.

25. A method as recited in claim 24, wherein each field indicates that a corresponding field of a packet received in the third processing device is to be used to generate an index to a table identifying the plurality of associated traffic handling systems, the generated index being associated with the selected traffic handling system.

5 26. A method as recited in claim 25, wherein each field indicates that a hashed value of the corresponding field of the received packet is to be used to generate the index to the table identifying the plurality of associated traffic handling systems.

27. A method as recited in claim 23, wherein at least one of the fields may be set to indicate one or more port identifiers of traffic received in the third processing device.

10 28. A method as recited in claim 27, wherein the fields are selected from a group consisting of a port 0 field, a port 1 field, a port 2 field, a port 3 field, a port 4 field, a port 5 field, a port 6 field, and a port 7 field.

15 29. A method as recited in claim 28, wherein the fields includes a source/destination field to indicate whether the port identifiers of the received traffic are source ports or destination ports.

30. A method as recited in claim 20, wherein the traffic information includes security options for specifying an authentication level for messages communicated between the third processing device and the selected traffic handling system.

20 31. A method as recited in claim 30, wherein the security options are configurable to select no authentication for messages communicated between the third processing device and the selected traffic handling system.

32. A method as recited in claim 30, wherein the security options are configurable to require a predetermined password encoded within messages communicated between the third processing device and the selected traffic handling system.

33. A method as recited in claim 20, wherein the traffic information includes identifying information for the selected traffic handling system.

34. A method of reinserting a redirected packet into a flow between a first processing device and a second processing device, the method comprising:

receiving a redirected packet from a third processing device in a fourth processing device; and

sending the packet back to the third processing device for transmission.

35. A method as recited in claim 34, further comprising encapsulating the redirected packet prior to sending the packet back so as to indicate that the packet is not to be redirected by the third processing device back to the fourth processing device.

36. A method as recited in claim 35, wherein the redirected packet is encapsulated in a GRE packet.

37. A method as recited in claim 34, further comprising modifying the redirected packet prior to sending the packet back.

38. A method as recited in claim 37, wherein a destination address of the redirected packet is modified prior to sending the packet back.

39. A method as recited in claim 34, further comprising making a copy of the redirected packet prior to sending the packet back.

40. A first computer system associated with a plurality of traffic handling systems and operable to facilitate redirection of traffic sent from a second computer system to a third computer system, the traffic being redirected to a selected traffic handling system, the computer system comprising:

a memory; and

a processor coupled to the memory,

wherein at least one of the memory and the processor are adapted to provide:

receiving traffic information from at least a portion of the associated traffic handling systems, wherein the traffic information specifies which data should be redirected to the portion of associated traffic handling systems;

communicating the traffic information to at least a designated one of the associated traffic handling systems; and

receiving traffic redirection information from the designated traffic handling system, the traffic redirection information being based on the communicated traffic information.

41. A computer system as recited in claim 40, wherein at least one of the memory and the processor are further adapted to provide:

building or updating a data structure based on the received traffic information, wherein the traffic information is communicated to the designated traffic handling system within the data structure.

42. A computer system as recited in claim 40, wherein at least one of the memory and the processor are further adapted to provide:

receiving a packet from a second computer system destined for a third computer system; and

5 redirecting the packet to a selected one of the traffic handling systems based on the traffic redirection information.

43. A computer system as recited in claim 42, wherein at least one of the memory and the processor are further adapted to provide:

receiving the packet back after redirecting it to the selected traffic handling system;

10 determining that the packet is to be sent to the packet's original destination address instead of being redirected to the selected traffic handling system; and

sending the packet to its original destination.

15 44. A computer system as recited in claim 43, wherein it is determined that the packet is to be sent to the packet's original destination by determining that the packet is encapsulated and de-encapsulating the packet prior to sending the packet to its original destination.

20 45. A computer system as recited in claim 40, wherein the traffic information sent from a selected traffic handling system includes service options specifying which data is to be redirected to the selected traffic handling system.

46. A computer system as recited in claim 45, wherein the service options include a plurality of fields that are configurable to indicate that one or more fields of a packet

CISCP146

received in the first computer system are to be used to determined redirection of packets to the selected traffic handling system.

47. A computer system as recited in claim 46, wherein the fields are selected from a group consisting of a source IP field, a destination IP field, a source port field, a destination port field, a source IP alternative field, a destination IP alternative field, a source port alternative field, and a destination port alternative field.

48. A computer system as recited in claim 47, wherein each field indicates that a corresponding field of a packet received in the first computer system is to be used to generate an index to a table identifying the plurality of associated traffic handling systems, the generated index being associated with the selected traffic handling system.

49. A computer system as recited in claim 48, wherein each field indicates that a hashed value of the corresponding field of the received packet is to be used to generate the index to the table identifying the plurality of associated traffic handling systems.

50. A computer system as recited in claim 46, wherein at least one of the fields may be set to indicate one or more port identifiers of traffic received in the computer system.

51. A computer system as recited in claim 50, wherein the fields are selected from a group consisting of a port 0 field, a port 1 field, a port 2 field, a port 3 field, a port 4 field, a port 5 field, a port 6 field, and a port 7 field.

52. A computer system as recited in claim 51, wherein the fields includes a source/destination field to indicate whether the port identifiers of the received traffic are source ports or destination ports.

53. A computer system as recited in claim 50, wherein at least one of the memory and the processor are further adapted to provide:

receiving a packet from a first processing device destined for a second processing device; and

when one or more port identifiers of the received packet matches a corresponding set field of the service options of the selected traffic handling system, redirecting the packet to the selected traffic handling system.

54. A computer system as recited in claim 40, wherein the traffic information sent from a selected traffic handling system to the first computer system includes security options for specifying an authentication level for messages communicated between the first computer system and the selected traffic handling system.

55. A computer system as recited in claim 54, wherein the security options are configurable to select no authentication for messages communicated between the first computer system and the selected traffic handling system.

56. A computer system as recited in claim 54, wherein the security options are configurable to require a predetermined password encoded within messages communicated between the first computer system and the selected traffic handling system.

57. A computer system as recited in claim 40, wherein the traffic information sent from a selected traffic handling system includes identifying information for the selected traffic handling system.

58. A designated traffic handling system associated with a plurality of traffic handling systems and operable to facilitate redirection of traffic sent from a first computer

system to a second computer system, the traffic being redirected by a third computer system to a selected traffic handling system, the designated traffic handling system comprising:

a memory; and

a processor coupled to the memory,

wherein at least one of the memory and the processor are adapted to provide:

receiving traffic information from the third computer system, wherein the traffic information specifies which data should be redirected to which associated traffic handling systems; and

allocating traffic portions to each traffic handling system based on the received traffic information.

59. A designated traffic handling system as recited in claim 58, wherein the traffic portions are allocated by:

building or updating a data structure based on the received traffic information within the third processing device that receives traffic sent from the first processing device to the second processing device, the data structure specifying how to direct traffic received in the third processing device to selected traffic handling systems associated with the third processing device.

60. A designated traffic handling system as recited in claim 58, wherein the traffic information includes service options specifying which data is to be redirected to the selected traffic handling system.

61. A designated traffic handling system as recited in claim 60, wherein the service options include a plurality of fields that are configurable to indicate that one or more

fields of a packet received in the third processing device are to be used to determined redirection of packets to the selected traffic handling system.

62. A designated traffic handling system as recited in claim 61, wherein the fields are selected from a group consisting of a source IP field, a destination IP field, a source port field, a destination port field, a source IP alternative field, a destination IP alternative field, a source port alternative field, and a destination port alternative field.

63. A designated traffic handling system as recited in claim 62, wherein each field indicates that a corresponding field of a packet received in the third processing device is to be used to generate an index to a table identifying the plurality of associated traffic handling systems, the generated index being associated with the selected traffic handling system.

64. A designated traffic handling system as recited in claim 63, wherein each field indicates that a hashed value of the corresponding field of the received packet is to be used to generate the index to the table identifying the plurality of associated traffic handling systems.

65. A method as recited in claim 61, wherein at least one of the fields may be set to indicate one or more port identifiers of traffic received in the third processing device.

66. A method as recited in claim 65, wherein the fields are selected from a group consisting of a port 0 field, a port 1 field, a port 2 field, a port 3 field, a port 4 field, a port 5 field, a port 6 field, and a port 7 field.

67. A designated traffic handling system as recited in claim 66, wherein the fields includes a source/destination field to indicate whether the port identifiers of the received traffic are source ports or destination ports.

68. A designated traffic handling system as recited in claim 58, wherein the traffic information includes security options for specifying an authentication level for messages communicated between the third processing device and the selected traffic handling system.

69. A designated traffic handling system as recited in claim 68, wherein the security options are configurable to select no authentication for messages communicated between the third processing device and the selected traffic handling system.

70. A designated traffic handling system as recited in claim 68, wherein the security options are configurable to require a predetermined password encoded within messages communicated between the third processing device and the selected traffic handling system.

71. A method as recited in claim 68, wherein the traffic information includes identifying information for the selected traffic handling system.

72. A first processing system operable to reinsert a redirected packet into a flow between a second processing device and a third processing device, the computer system comprising:

- a memory; and
- a processor coupled to the memory,

wherein at least one of the memory and the processor are adapted to provide:

receiving a redirected packet from a fourth processing device in the first processing device; and

5 sending the packet back to the fourth processing device for transmission.

73. A processing system as recited in claim 72, wherein at least one of the memory and the processor are further adapted to provide encapsulating the redirected packet prior to sending the packet back so as to indicate that the packet is not to be redirected by the fourth processing device back to the first processing device.

74. A processing system as recited in claim 73, wherein the redirected packet is encapsulated in a GRE packet.

75. A processing system as recited in claim 72, wherein at least one of the memory and the processor are further adapted to provide modifying the redirected packet prior to sending the packet back.

76. A processing system as recited in claim 75, wherein a destination address of the redirected packet is modified prior to sending the packet back.

77. A processing system as recited in claim 72, wherein at least one of the memory and the processor are further adapted to provide making a copy of the redirected packet prior to sending the packet back.

78. A computer program product for facilitating redirection of traffic sent from a first processing device to a second processing device, the computer program product comprising:

at least one computer readable medium;

computer program instructions stored within the at least one computer readable product configured to cause a processing device to provide:

at a third processing device associated with a plurality of traffic handling systems, receiving traffic information from at least a portion of the associated traffic handling systems, wherein the traffic information specifies which data should be redirected to the portion of associated traffic handling systems;

communicating the traffic information to at least a designated one of the associated traffic handling systems; and

at the third processing device, receiving traffic redirection information from the designated traffic handling system, the traffic redirection information being based on the communicated traffic information.

79. A computer program product for facilitating redirection of traffic sent from a first processing device to a second processing device, the computer program product comprising:

at least one computer readable medium;

computer program instructions stored within the at least one computer readable product configured to cause a processing device to provide:

at a designated traffic handling system associated with a plurality of traffic handling systems and one or more routers, receiving traffic

METHOD AND APPARATUS FOR REDIRECTING NETWORK TRAFFIC

ABSTRACT OF THE DISCLOSURE

Methods and apparatus are described for intelligently redirecting and/or processing traffic. In one embodiment, a method of facilitating redirection of traffic sent from a first processing device (e.g., a client) to a second processing device (e.g., a server) is disclosed. At a third processing device (e.g., a router) associated with a plurality of traffic handling systems (e.g., cache systems), traffic information is received from at least a portion of the associated traffic handling systems. The traffic information specifies which data should be redirected to the portion of associated traffic handling systems. It is determined how to redirect data received by the third processing device to a selected traffic handling system based on the received traffic information received from at least one of the associated traffic handling systems. In one aspect, the determination of redirecting data is accomplished by communicating the traffic information to at least a designated one of the associated traffic handling systems and receiving traffic redirection information from the designated traffic handling system at the third processing device. In another aspect, a packet is received in the third processing device from the first processing device destined for the second processing device, and the packet is redirected to a selected one of the traffic handling systems based on the traffic redirection information. In another method implementation, at a designated traffic handling system associated with a plurality of traffic handling systems and one or more routers, traffic information is received from at least one of the routers. The traffic information specifies which data should be redirected to which associated traffic handling systems. Traffic portions are allocated to each traffic handling system based on the received traffic information. In another implementation, each traffic handling system may determine that particular redirected packets should be transmitted back to the redirecting router and reinserted into the original traffic flow.




Fig. 1

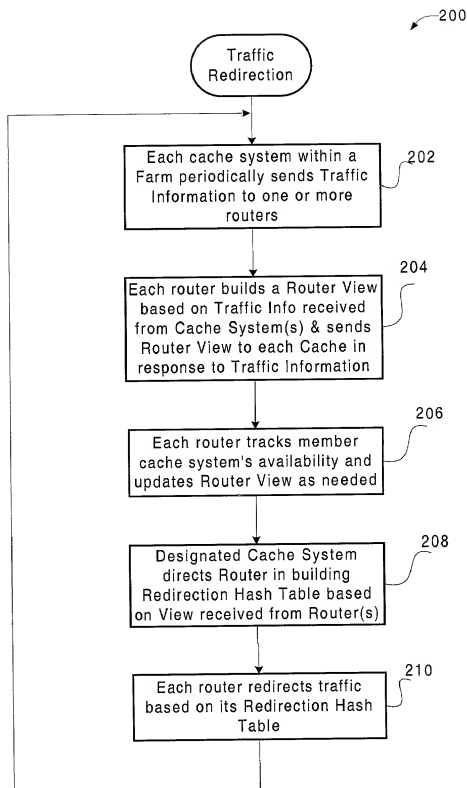


Fig. 2

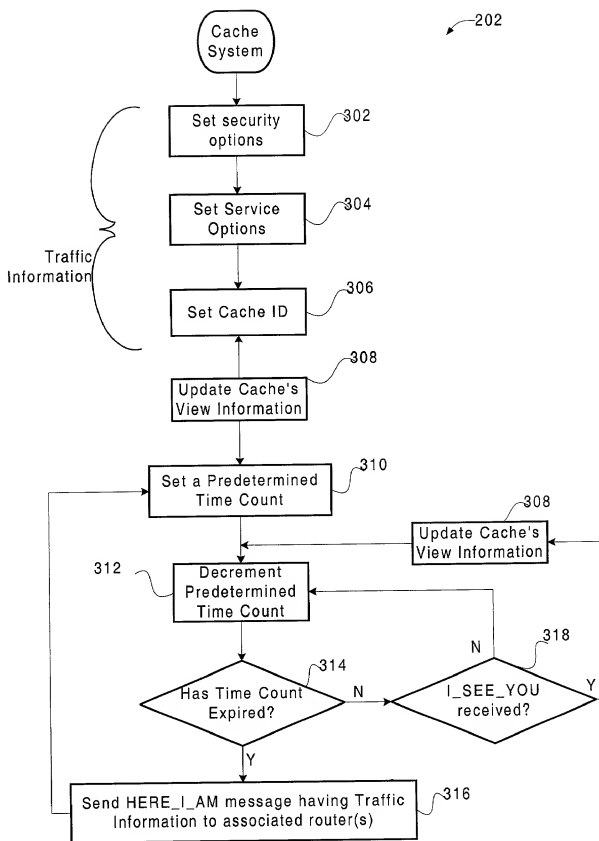


Fig. 3

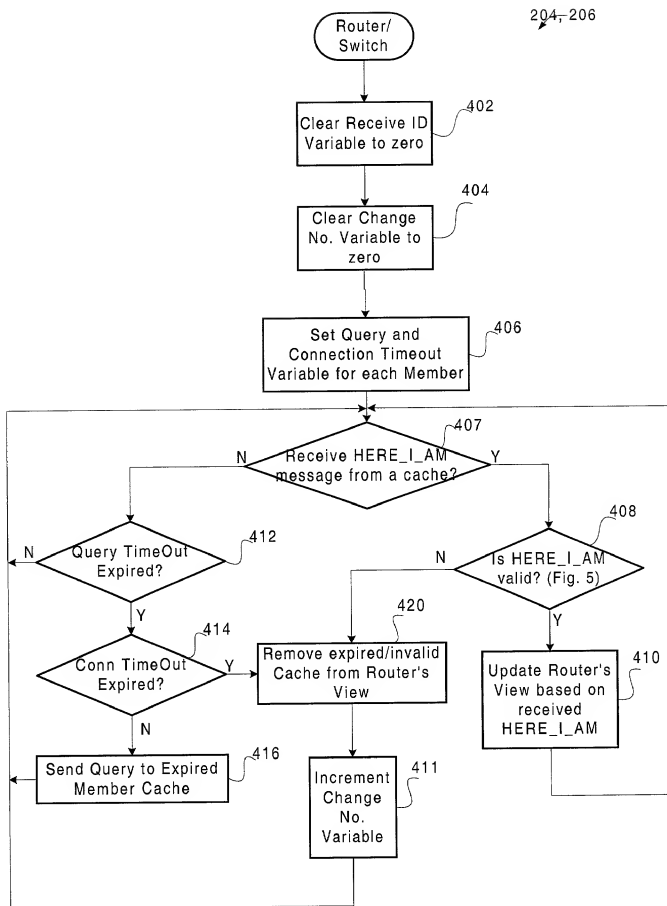


Fig. 4

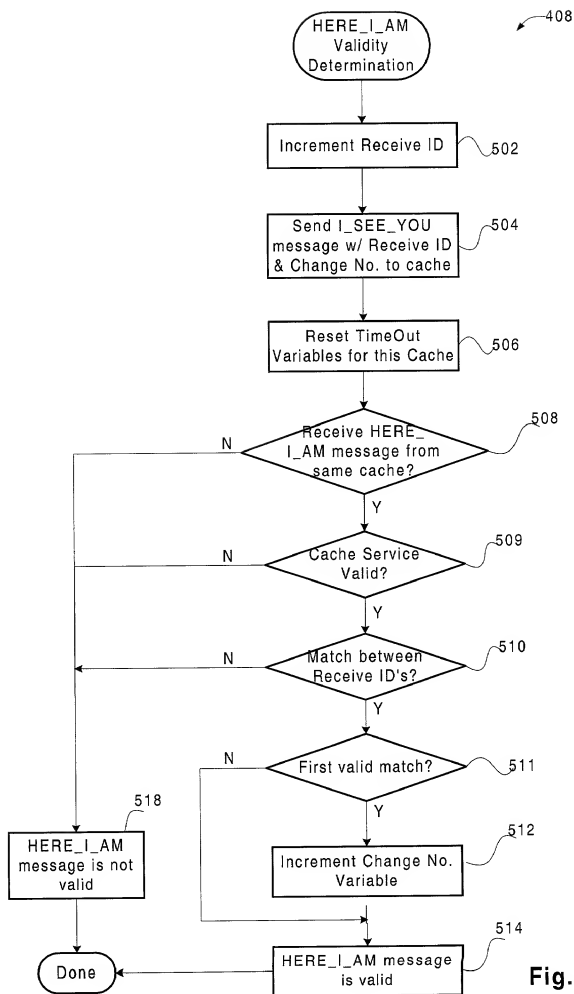


Fig. 5

